



Defense Science Board
Report On
Advanced Computing

March 2009

*Office of the Under Secretary of
Defense for Acquisition, Technology
and Logistics*

Washington, D.C. 20301-3140

| Report Documentation Page | | | | Form Approved OMB No. 0704-0188 | |
|--|------------------------------------|-------------------------------------|---|---|---------------------------------|
| Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. | | | | | |
| 1. REPORT DATE MAR 2009 | | 2. REPORT TYPE | | 3. DATES COVERED 00-00-2009 to 00-00-2009 | |
| 4. TITLE AND SUBTITLE Advanced Computing | | | | 5a. CONTRACT NUMBER | |
| | | | | 5b. GRANT NUMBER | |
| | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) | | | | 5d. PROJECT NUMBER | |
| | | | | 5e. TASK NUMBER | |
| | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Office of the Under Secretary of Defense for ,Acquisition, Technology and Logistics,Washington,DC,20301-3140 | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) | |
| | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | | | |
| 13. SUPPLEMENTARY NOTES | | | | | |
| 14. ABSTRACT | | | | | |
| 15. SUBJECT TERMS | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT Same as Report (SAR) | 18. NUMBER OF PAGES 64 | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | | | |

This report is a product of the Defense Science Board (DSB).

The DSB is a Federal Advisory Committee established to provide independent advice to the Secretary of Defense. Statements, opinions, conclusions, and recommendations in this report do not necessarily represent the official position of the Department of Defense.

This report is UNCLASSIFIED and is releasable to the public.



DEFENSE SCIENCE
BOARD

OFFICE OF THE SECRETARY OF DEFENSE
3140 Defense Pentagon
Washington, DC 20301-3140

17 February 2009

MEMORANDUM FOR UNDER SECRETARY OF DEFENSE FOR ACQUISITION, TECHNOLOGY AND LOGISTICS

SUBJECT: The Final Report of the Defense Science Board (DSB) Task Force on the National Nuclear Security Administration's (NNSA) Strategic Plan for Advanced Computing

I am pleased to forward to you the final report of the DSB Task Force on NNSA's Strategic Plan for Advanced Computing, co-chaired by Dr. Bruce Tarter and Mr. Robert Nesbit.

The Task Force was asked to evaluate NNSA's strategic plan for Advanced Simulation and Computing (ASC) and its adequacy to support the Stockpile Stewardship Program (SSP), whose mission is to ensure the safety, performance and reliability of our Nation's nuclear weapons stockpile. The Task Force was also asked to evaluate the role of ASC in maintaining US leadership in advanced computing and assess the impact of using ASC's capabilities for broader national security and other issues.

The Task Force concluded that, since the cessation of nuclear testing, ASC has taken on the principal integrating role in assuring the long term safety and reliability of the stockpile. It is also an essential tool in addressing specific stockpile issues. Furthermore, ASC has played a leadership role in re-establishing US leadership in high performance computing. The use of ASC and ASC-derived technology for other national security, scientific, and commercial applications has also increased dramatically, and high performance computing is viewed as an extremely valuable and cost-effective approach to many of the user's important problems.

However, it is not likely that ASC will meet the compelling goals stated in its roadmaps and planning documents at the currently projected levels of funding. Furthermore, the high end of the US computing industry may be negatively impacted with implications for the much broader range of potential users in the DOD, other federal agencies, and the commercial world. Accordingly, the Task Force strongly recommends sizing the budget of ASC to meet its nuclear weapons objectives and retain US leadership in advanced computing.

I fully endorse all of the Task Force's recommendations and urge you to review this report and give special consideration to their findings and recommendations.

A handwritten signature in black ink, reading "William Schneider, Jr.", is positioned above the printed name.

Dr. William Schneider, Jr.
DSB Chairman

This Page is Intentionally Left Blank



DEFENSE SCIENCE
BOARD

OFFICE OF THE SECRETARY OF DEFENSE
3140 Defense Pentagon
Washington, DC 20301-3140

16 February 2009

MEMORANDUM FOR THE CHAIRMAN OF THE DEFENSE SCIENCE BOARD

SUBJECT: Final Report of the Defense Science Board (DSB) Task Force on the National Nuclear Security Administration (NNSA) Strategic Plan for Advanced Computing

We are pleased to present to you our final report which describes our assessment of NNSA's strategic plan for Advanced Simulation and Computing (ASC). As requested in the Terms of Reference we have also evaluated the impact of ASC in maintaining US leadership in high performance computing (HPC) and of using the planned HPC capabilities for broader national security and other issues.

To carry out this study, the Task Force held five meetings between April and October of 2008, during which time we received more than 70 briefings from: NNSA representatives involved in HPC; scientists from the NNSA Labs in Livermore, Los Alamos and Sandia; representatives from most of the Federal Agencies involved in HPC, and from individuals leading the work in HPC for the major industrial users. We also reviewed numerous planning documents provided by the NNSA/ASC program although no formal strategic plan exists and no resource numbers were attached to the plans (which prevented us from meeting the precise letter of our terms of reference).

In brief the Task Force concluded that:

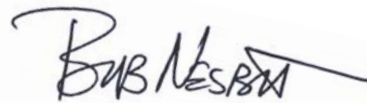
1. Since the cessation of nuclear testing ASC has taken on the primary integrating role in assuring the safety and reliability of the Nation's nuclear stockpile. It is the principal tool in combining nuclear test history, data from laboratory experiments, and weapons designer expertise into an improved understanding of weapon performance, reliability, safety, and security. It has provided the means to resolve significant issues with the nuclear weapons stockpile.
2. ASC and its ASCI predecessor program have played a leadership role in regaining US leadership in HPC, and ASC computers occupy the top rungs of the world list of most powerful computers.
3. ASC has significantly contributed to the advancement of high performance computing technologies widely used by other federal agencies and some commercial sectors. There are a number of application areas where HPC plays an increasing role: national security (e.g. in nuclear forensics); energy and environmental science (e.g. global climate); and the commercial world (e.g. exploration for natural resources).

4. The ASC program needs significantly more resources in the future to achieve the goals stated in its roadmaps and planning documents. At currently projected levels of funding it will not meet its nuclear weapons milestones in a timely manner and perhaps not at all. Thus the goal of a predictive capability for nuclear weapons design, which many feel is essential for making significant modifications to the stockpile, is unlikely to be achieved with present program plans and projected resource levels.
5. The development of the next levels of HPC, i.e. computational capability in the many petaflop and possibly the exaflop regime, will be significantly more challenging than the already difficult climb to the current level (approaching one petaflop for practical problems). Thus, it will require proportionately more resources to have a realistic chance of reaching these performance levels.

We are very appreciative of the time and effort put forth by the leadership of NNSA's Office of Research and Development for National Security, Science and Technology; by the laboratory staff who hosted the Task Force at Livermore, Los Alamos and Sandia; and are especially grateful for all of the federal agency and industry representatives who helped inform the Task Force members on this most important issue.

A handwritten signature in dark ink, appearing to read 'Bruce Tarter', with a stylized, flowing script.

Dr. Bruce Tarter
Co-Chairman

A handwritten signature in dark ink, appearing to read 'Bob Nesbit', with a stylized, flowing script.

Mr. Robert Nesbit
Co-Chairman

Table of Contents

| | |
|---|----|
| Findings and Recommendations | 3 |
| Introduction | 7 |
| The Role of High Performance Computing | 11 |
| Other DOE and National Security Missions | 21 |
| The Role of ASC/HPC in the Work of Other Organizations..... | 25 |
| Computer Matters | 33 |
| References | 43 |
| Appendix A: Terms of Reference..... | 45 |
| Appendix B: Task Force Membership..... | 47 |
| Appendix C: List of Briefings..... | 51 |
| Appendix D: Acronyms and Initialisms | 55 |

This Page is Intentionally Left Blank

Findings and Recommendations

The Defense Science Board Task Force on the National Nuclear Security Administration (NNSA) Strategic Plan for Advanced Computing was asked in its Terms of Reference (see Appendix A) to assess a number of topics, which can be summarized as follows:

- The adequacy of the NNSA's strategic plan for high performance computing (HPC) in supporting the Stockpile Stewardship Program.
- The role of, and the impacts of changes in investment on, research and development of high-performance computing supported by the NNSA in fulfilling its mission and maintaining the leadership of the United States in high performance computing.
- The importance of using current and projected scientific computing capabilities of the NNSA and other agencies to address a broad spectrum of national security challenges.
- The efforts of the Department of Energy to coordinate and develop joint strategies within its own department, with other agencies, and with the commercial sector to develop and apply high performance computing capabilities.

To carry out this assessment, the Task Force held five meetings between April and October of 2008, three in Washington, D. C. and two at the NNSA Laboratories in California and New Mexico. Briefings were delivered by NNSA representatives involved in HPC, scientists from the NNSA Labs in Livermore, Los Alamos and Sandia, representatives from most of the federal agencies involved in HPC, and by individuals leading HPC work for major industrial users.

The Task Force's major findings and recommendations are as follows:

Findings

- High performance computing (HPC) has been a principal nuclear design tool since the beginning of the nuclear weapons program. Following the cessation of nuclear testing in 1992, HPC has taken on the primary integrating role in assuring the safety and reliability of the stockpile.
- NNSA's Advanced Simulation and Computing (ASC) program and its predecessor, the Accelerated Strategic Computing Initiative (ASCI), have provided the means to combine nuclear test history, data from laboratory experiments and weapons designer expertise into a significantly improved understanding of nuclear

weapon performance, reliability, safety and security. This has led to a number of examples in which HPC has been a central element in stockpile stewardship decision making, e.g. whether observed stockpile issues would require major (and expensive) stockpile refurbishment.

- ASC budgets have declined significantly since FY02. The average yearly decrease has been between 5 and 10% depending on what factor is chosen for inflation, and workforce levels devoted to weapons computing have decreased by approximately half. Future budgets are projected as flat or declining.
- There are a number of key unresolved issues in our understanding of nuclear weapons. No formal strategic plan for advanced computing exists at NNSA. However, ASC has a reasonable roadmap with a set of well-defined milestones over the next several years to develop and acquire the next generation of high performance computing capability to attack these issues. If, as NNSA/ASC officials often stated, the likely budget scenario for ASC is one of flat or declining budgets (before inflation), then it is impossible to follow the ASC roadmap without compromising its goals and/or timescale. Future program needs cannot be met in a timely way at the projected resource levels. These programs needs include: full three dimensional (3D) simulations to address significant findings (SFIs) in an aging stockpile; potential stockpile modifications that move increasingly farther away from the legacy stockpile; and the incorporation into the stewardship program of results from the new experimental facilities such as the National Ignition Facility (NIF) and the Dual Axis Radiographic Hydrodynamic Test (DARHT) facility.

If, as NNSA/ASC officials often stated, the most optimistic future budget scenario is one of flat budgets (before inflation), then it is impossible to follow the ASC roadmap without compromising its goals and/or timescale

- The projected reduced ASC budgets are also inadequate to support strong peer review among the design laboratories including the development and maintenance of different computational approaches. A single computational method, code, or team is not a move toward efficiency; rather, it is a recipe for single point failure. Failing to follow through on the ASC plans will introduce considerable future risk into the nuclear weapons program.
- The Task Force found widespread use of HPC in other Federal Agencies and certain sectors of the commercial world (albeit often somewhat behind the state of the art in ASC). ASC was of great implicit benefit to these organizations either through their use of the new commercially available computers or through the custom modifications of technologies which ASC helped create. There is also general recognition of the leadership role played by NNSA/ASC in pushing the

state of the art in computing capability through their partnerships with multiple vendors.

- The Secretary of Energy and NNSA Administrator have called for broadening the support base for leading edge HPC both within DOE and by other agencies. However, even within appropriate program areas under their jurisdiction they have not yet made programmatic and funding commitments to make such broadening occur. We strongly encourage the new Administration to take such actions within DOE/NNSA (the partnership with the Office of Science is admirable and effective but has existed for some time and does not represent a broadening of the base).
- The Task Force has identified two potential security issues based on our understanding of NNSA-ASC's desire to share computing resources among different classification levels. The first security issue is the idea of "swinging" a machine between classified and unclassified uses, which has the potential of exposing a classified machine to the internet. The second, more subtle, security issue has to do with using the machine for different types of classified applications with different levels of classification. While multi-level security has been a long term goal, it is not yet a reality. Although the NNSA community is very cognizant of the sensitivity of nuclear weapons information, only a small fraction have worked with intelligence-related data which has a quite different set of sensitivities concerning handling and distribution of data.
- The computer and computational science plans for the next half decade (out to 10s of petaflop machines) are challenging but probably within the reach of the industry and applications communities. The following generation of computers will require extensive research and development to have a chance of reaching the exascale level. Even if exascale level machines can be created there are extremely difficult challenges in their use for core NNSA applications.

Recommendations

- ASC should develop and frequently update a formal strategic plan. It should combine the elements of its other planning documents and include projected resource levels.
- ASC budgets should be sized to provide adequate funding for the computer development and programmatic applications needed to meet the stated goals of the nuclear weapons program. The level of resources should be sufficient to ensure that critical work force levels are maintained, that multiple approaches to complex computational issues are pursued, and that several vendors remain at the leading edge of supercomputing capability in the U.S. In addition, the ASC program needs to be organized to analyze and exploit the capabilities of the new SSP experimental facilities (such as NIF and DARHT) and to translate their results into weapons impacts. This will be required whether the emphasis is on maintaining the legacy stockpile or making more significant modifications to the future stockpile.

ASC budgets should be sized to provide adequate funding for the computer development and programmatic applications needed to meet the stated goals of the nuclear weapons program

- The Task Force recommends aggressively pursuing the ASC program to help assure that HPC advances are available to the broad national security community. As in the past, many other national security organizations will use the ASC developed capabilities for their own needs. The DOE should enhance its own efforts by further strengthening the partnership between the NNSA and the Office of Science and then developing an HPC element in its other mission areas (such as Nuclear Non-Proliferation and Nuclear Energy).
- NNSA should seek the views of experts in cyber security before expanding into some of the potential uses of NNSA classified machines. While it is notable that Sandia has devoted considerable effort to creating safe mechanisms for sharing machines, there has always been a balance between the laudable efficiency goals and the current threat profile. It is time for a re-examination of the issue.
- The Task Force recommends including a significant level of research and development funds in its pursuit of the next generations of petascale and then exascale level computing capability. This includes both the hardware and the complex software that may be required for the architectures needed for exascale capability. The challenges are extremely daunting, especially at the exascale level. Only a broadly based effort including multiple approaches to the hardest problems is likely to produce success for the ASC/NNSA mission and maintain U.S. leadership in HPC.

Introduction

The Stockpile Stewardship Program focuses on surveying, assessing and refurbishing the stockpile without the need for nuclear testing

During the Cold War, nuclear weapons entered the stockpile through a design, test and build sequence. The stockpiled weapons were periodically evaluated, altered, and eventually retired. A new warhead type was introduced into the stockpile (i.e., carried through the design, testing and production sequence) every year or two, and there were generally several nuclear weapons in the “pipeline” at any one time. New nuclear warheads were designed in direct response to military requirements and/or were driven by technological possibilities that were then adopted by the military. These new nuclear explosive designs were simulated in great detail using high performance computers and laboratory-scale experiments, and then tested in integral full-scale nuclear explosive experiments. Once a design type was accepted by the military, typically after a competition between the two design laboratories, it was engineered for the intended application and manufactured by the selected production complex, which received and assembled components provided by various sites. The weapons in the stockpile were surveilled, assessed (sometimes with nuclear tests) and occasionally refurbished, but the program was dominated by the frequent introduction of new designs and the retirement of old ones.

Nuclear testing and new warhead design and production ceased altogether following the end of the Cold War. The last U.S. nuclear weapon test was on September 23, 1992, and no new designs have been introduced into the stockpile since the W88 in 1989.

In 1993, the Stockpile Stewardship Program (SSP) was created with a goal of maintaining the safety and reliability of the existing stockpile without the need for nuclear testing. This program became the centerpiece of the nuclear weapons program following the signing (although not the ratification) of the Comprehensive Test Ban Treaty (CTBT) in 1996. The SSP was founded on the belief that these goals could be achieved by preserving and reinvigorating the intellectual base of the Laboratories; employing an array of advanced computers, modeling approaches, and experimental techniques; and implementing a more comprehensive stockpile surveillance and refurbishment program.

The SSP replaced the design-test-build sequence of the Cold War with a sequence focused on surveying, assessing and refurbishing the stockpile, coupled with a vigorous scientific program to gain a better understanding of nuclear weapons in the absence of nuclear testing. Any issues found during

the surveillance process (e.g., aging problems such as cracks or corrosion) are assessed for their impact on the safety and performance of the weapon using a family of advanced supercomputer codes and new laboratory facilities. Problems are then corrected by refurbishment of the warhead using the production complex. Furthermore, a schedule of systematic maintenance and upgrading was instituted. In this Life Extension Program (LEP), each warhead type is refurbished on a scheduled basis to ensure the long-term health of the stockpile and more cost-efficient workload balancing within the complex.

A major part of the SSP is an effort to better understand the science involved in nuclear explosions. The objective is to reduce uncertainties so that the level

of confidence in an assessment of weapon performance and safety is comparable with that once achieved through a combination of computer calculations, non-nuclear experiments and nuclear tests, but now without nuclear tests. Ultimately, this led to the development of the Quantification of Margins and Uncertainties (QMU) approach, which is a systematic way of evaluating the performance margin of a nuclear warhead. As long as the margin is large compared with the technical uncertainties, there should be confidence in the nuclear performance of the warhead.

More than a decade after its inception, the SSP has accumulated a body of substantial achievements. The program has made significant advances in the basic science of nuclear weapons performance and the properties of nuclear explosive materials. It has led to the development and certification of new processes for manufacturing plutonium pits, as well as the establishment of a systematic process that is vetted and applied on an annual basis to certify the U.S. nuclear stockpile. These achievements were possible because the SSP challenged and rejuvenated the technical personnel in each of the Laboratories associated with the nuclear weapons program by supplying them with the resources and facilities they needed to do their new job. In particular, SSP built the world's greatest supercomputing capability and applied it successfully in the ASCI and ASC programs to understand and help mitigate stockpile issues. It has constructed, or is in the process of constructing, state-of-the-art laboratory facilities, including: the National Ignition Facility (NIF) at Lawrence Livermore National Laboratory (LLNL), which will help advance understanding of material properties at nuclear weapon conditions not previously achievable in the Laboratory; the Dual Axis Radiographic Hydrodynamic Test facility (DARHT) at Los Alamos National

Quantification of Margins and Uncertainties (QMU) is a systematic way of evaluating the performance margin of the nuclear warhead. Only through modeling and simulation can we demonstrate the safety margins of a particular warhead

Laboratory (LANL), which creates intense bursts of X-rays that are used to create digital images of mock nuclear devices as they implode; the Z machine at Sandia National Laboratory (SNL), which is designed to study fusion; and a sub-critical experiments capability at the Nevada Test Site (NTS). These facilities will provide new insights into weapons science and weapon performance. The SSP has used these new computational and experimental tools to resolve many issues from earlier tests and to teach a new generation of scientists about the stockpile and nuclear design.

SSP built the world's greatest supercomputing capability and applied it successfully to the ASCI and ASC programs

However, concerns remain for the long term maintenance of the Cold War stockpile (often referred to as the legacy stockpile), as well as its applicability to future deterrence in a more pluralistic world. To this end, the concept of the Reliable Replacement Warhead (RRW) program was introduced as a means to upgrade the legacy stockpile by replacing one or more of its nine systems by militarily equivalent but technologically more robust warheads. These warheads would be developed using the extensive test data base and high performance computers and entered into the stockpile without a new nuclear test. The first RRW design competition, held between 2005-2006, aimed to develop a replacement for some of the Submarine Launched Ballistic Missiles (SLBMs). While the project was awarded to LLNL, the program is on hold pending Congressional approval, and satisfactory resolution of Congressional questions regarding the Nation's overall nuclear weapons posture.

No matter how that discussion turns out, there is a reasonable consensus that any future model of the nuclear weapons complex must include a modernized industrial base that can refurbish or make weapons at a lower cost than at present, and in a more efficient, safer, and environmentally benign manner. To accomplish this goal the NNSA has proposed a Complex Transformation Plan which would substantially upgrade or rebuild major elements of the production system while reducing operations at a number of sites. These plans, while not costed on any comprehensive basis, will certainly require significant initial investments for a period of at least several years. At the same time, the major experimental stewardship facilities, such as the NIF and DARHT are just now coming on line, and in conjunction with the next generation of high performance computers, will need to produce an extended body of work to meet the objectives of the SSP. If, as is often stated by NNSA officials, their most optimistic budgetary scenario is one of flat budgets, (before inflation) then it is impossible to fit all of these plans into the overall nuclear weapons program without compromising either its goals, timescale, or both. Although comments on NNSA's overall priorities are beyond the scope of this study, we can address the implications that


reduced budgets and stretched out time horizons can have on ASC, and by implication, other elements of the SSP. That is the background against which the discussion of ASC takes place in this report.

The Role of High Performance Computing

Background

High performance computing has been an essential core ingredient of the nuclear weapons program since its inception in the early 1940s. From the Electronic Numerical Integrator and Computer (ENIAC) in World War II

(WWII) to the present ASC machines, most of the actual nuclear weapons design and testing has been done on the most advanced electronic computers available at any given time. In many cases, the development of the next generation of such computers is done at the request of and in tandem with the nuclear weapons community.



U.S. Army Photo: The ENIAC, in BRL building 328. Left: Glen Beck Right: Frances Elizabeth Snyder Holberton

So what does it mean to “design” a nuclear weapon? Like any design activity it starts with a diagram or sketch of where all the parts go and how they connect together.

For a nuclear weapon, the parts

include the plutonium or uranium, the high explosive, the firing system, safety devices, the delivery vehicle it has to fit into, and all of the interconnecting pieces that have to remain functional for years in a high radiation, chemically reactive environment created by the materials used in constructing the bomb. For the modern warheads, it also involves the heavy isotopes of hydrogen—deuterium and tritium—that boost the primary yield and fuel the thermonuclear stages that greatly increase the overall yield. The art of weapon design consists of arranging these constituents in such a way as to maximize the yield to weight (for the historical stockpile), the performance margin for a successful explosion, the operational safety and security elements, and other features while minimizing the cost and difficulty of manufacturing the weapon.

Once the designer has an initial proposed configuration of the warhead, i.e. the amount and arrangement of all the parts as they will be constructed, the issue is how well it will work and meet the objectives of the military customer. To evaluate this question, the designer performs a series of numerical experiments by modeling the performance of the device on a

computer. After telling the computer the initial layout, the designer starts the calculation by “lighting” the fuse—just as on any explosive—and watching how the explosion develops. The computer models the process by solving the equations of motion and energy for all parts of the warhead, sequentially in time, until the explosion is complete. At each step, the computer has to have knowledge in every part of the device of the temperature, density, pressure, what chemical and nuclear reactions are occurring, how strong or brittle all of the materials are, and what happens when the bomb components mix together during various phases of the explosion.

The modeling process is a simplified version of what happens during the actual explosion. For example, the models often assume greater symmetry than is actually true (e.g. that an initial spherical configuration remains spherical since it is difficult and time-consuming to calculate the misalignment). Similarly, materials may be kept artificially homogeneous over large regions of an explosive; or chemical and physical properties are described by simple formulas in all regions. The designer runs a broad spectrum of numerical simulations to see which of these approximations matter and which are unimportant. For example, the compressibility of some material might be numerically increased by 20% compared to its assumed value to see how that affects the answer; or the amount of plutonium decreased by 5% and so on to see where the explosion’s regions of sensitivity are greatest. Substantial skill is needed to determine stable regions where small variations in construction or operating environment will minimally affect the actual performance of the device.

In addition to all of the computer experiments, the designer often requests special measurements by chemists, physicists or engineers to improve the data on important parts of the explosive. When the designer believes a satisfactory configuration has been reached, there is usually a full scale calculation of the entire explosion carried out from beginning to end with as much detail as can be put into the problem. Such computations typically take ten to one hundred hours to run on the largest supercomputer and might have to be done over several weekends or even months of actual time.

Prior to the end of the Cold War, the next step would then be to assemble the explosive and set it off at the Pacific Proving Grounds (until the early 1960s) or at the Nevada Test Site. The measurements would then establish how well the designer had predicted the yield, the output of various kinds of radiation, and the timing of various phases of the device. Because the explosion happens so quickly and under such extreme conditions, the diagnostic instruments usually measure only a small fraction of the information needed to understand the details of the explosive’s actual behavior. However, the designer (and everyone else) usually has a general

idea of how well things worked. Equally important, the live test allows the designer to calibrate the uncertainties in the computer models, and over time, establish various semi-empirical ways of treating the uncertainties in the simulation. As this process is repeated for different classes of explosives, and by different designers, the semi-empirical factors become codified as “computational knobs” that are used in simulations to bring the results into closer agreement with the measured test results (and to better predict the behavior of future designs).

Subsequent to the end of the Cold War, and the cessation of nuclear testing, the “designer’s” job changed significantly. Instead of developing new weapons, their task now was to steward an existing stockpile into the indefinite future. Detailed simulations are no longer of hypothetical weapons. They are carried out on existing weapons where aging or other issues arise and potential flaws are discovered in operating the weapon in a particular environment.

Among the most significant changes, however, is in the kind of calculations needed. In designing a new weapon, the configuration is under the designer’s control and there is usually a great deal of symmetry involved. Many one dimensional computations are done for sensitivity studies and the full weapons calculation often involves only two dimensional computer codes. In contrast, for an existing stockpile, a weapon (like a human or a car) ages in three dimensions; corrosion or a crack occurs on one side of a device, not equally on both sides. That means a numerical analysis requires a three dimensional computer code, and since each dimension is typically described by a thousand or so “grid points,” this means a thousand times more calculations are required. Next, since cracks and corrosion are initially small features, the resolution has to improve by approximately a factor of 10 (in order to “see” the crack), and another factor of a few to input the chemistry or physics of the aging process. Overall, a stewardship computer must be something like 10,000- 100,000 times more powerful than its predecessor to do its assigned tasks. This analysis set the requirements for the initial ASCI program computers.


Subsequently, the last decade has seen the development of a remarkable set of computers and three dimensional codes, with extraordinary graphics, that enable weapons scientists to probe areas of science and weapon behavior never before possible (and now required). Scientists have successfully utilized the new codes to carry out LEPs on several systems and to resolve several important Significant Findings (SFIs). And, not accidentally, the measured increase in performance from the beginning of stewardship until the ASCI-Purple machine in the present day shows about a factor of 10,000 in supercomputer performance.

Future Program Requirements

The National Ignition Facility



Experiments conducted on NIF will make significant contributions to nuclear weapons science. It will lead to major advances in three areas: understanding of material properties at nuclear weapon conditions not previously achievable in the Laboratory; resolution of major unsolved weapons problems in energy transport and thermonuclear burn; and validation of the advanced computer codes being developed to provide predictive capability for the stockpile.



NIF's ability to do experiments with complex targets under controlled conditions will be a primary tool in assessing a wide range of areas in which SFI's are likely to occur in the future.

However, as described in many official DOE-NNSA publications, as well as articles and testimony from Lab scientists, the future evolution of the nuclear weapons program appears likely to be much more complicated. Although there is no strong consensus on the size and diversity of the future stockpile, there is close to unanimity about the need for a modernized industrial complex for weapons and for a plan to refurbish/replace a significant fraction of the existing stockpile. This may take the form of aggressive LEPs or some form of replacement warheads, but in both cases there is a clear requirement for simulations that can confidently predict a weapons behavior farther away from the baseline configurations of the legacy weapons.¹ Meanwhile, there is a simultaneous need to use the new experimental stewardship machines, such as NIF and DARHT, to test new codes beyond the legacy nuclear test data. The net result is a need for an increase of at least a factor of 100 in computer capability, and perhaps considerably more to respond to the long term needs of a nuclear weapons program that must make substantial technical modifications to the existing stockpile without nuclear testing. That is the conclusion that drives the path forward for the Advanced Simulation and Computing Program.

The Dual Axis Radiographic Hydrodynamic Testing Facility (DARHT)



DARHT consists of two electron accelerators positioned at a 90-degree angle, each focused on a single firing point. It is at this point where nuclear weapon mock-ups are driven to extreme temperatures and pressures with high explosives and where the DARHT electron beams produce high-energy X-rays used to image the behavior of materials and systems under those extreme conditions. DARHT is a tool used to ensure the integrity of the nation's nuclear stockpile without nuclear testing.

technical goals of the program matched with the intended computer hardware and software.

The PCF laid out pegposts in six areas: Safety and Surety; Nuclear Explosive Package Assessment; Output, Effects and Survivability; Engineering Assessment; QMU and Validation and verification (V&V); and Experimental & Computational Capabilities. The Task Force heard presentations (frequently at the Secret Restricted Data (SRD) level) on many of these topics from both

The presumed Strategic Plan for ASC comprises a compilation of documents that address various aspects of the program. There is a ten year perspective on ASC, an ASC Business Model, a Platform Strategy, and an ASC Roadmap.^{2, 3, 4} Each of these addresses various aspects of future planning (with a good deal of self-consistent overlap) but without detailed resource numbers. Their integration is captured in a Predicative Capability Framework (PCF), in which milestones are delineated for the next decade in a half dozen areas important to the nuclear weapons program. Missing from all of these documents, however, is specificity of the resources needed, or to be allocated, to achieve the stated milestones. Consequently, it is difficult to assess the presumed plan's adequacy in the absence of such information. The Task Force did receive draft resource plans and scenarios, and the report will return to these after commenting on the



DARHT's electron accelerators use large, circular aluminum structures to create magnetic fields that focus and steer a stream of electrons down the length of the accelerator. Tremendous electrical energy is added along the way. When the stream of high-speed electrons exits the accelerator it is "stopped" by a tungsten target resulting in an intense burst of X-rays that are used to create digital images of mock nuclear devices as they implode.

DOE-NNSA and Laboratory staff, and received, or were referred to, a number of related documents and reports. Simultaneously, the Task Force heard about the general increase in high performance computing capability that is needed to reach these pegposts, and in nearly all cases, there exists a forceful set of arguments that the necessary level of two dimensional (2D) or 3D, full physics, and high resolution simulations will require computer capability that extends well into the petaflop regime and conceivably up to exaflops. There is also a view that the next generation of weapons workhorse computers could probably be developed and deployed for these tasks, but that the following generation of computers will face much more formidable issues in both hardware and software.

A better understanding of weapon performance will allow the inclusion of more stringent safety and security features without reducing confidence in device performance

Successful attainment of the pegposts in each area could have a significant impact on future LEPs or the design of replacement warheads, and by implication, on the costs associated with such efforts. For example, a better understanding of weapon performance might allow the inclusion of much more stringent safety or security features without reducing confidence in device

performance. Also, improved calculations of energy balance could give the designer the freedom to use different materials that lower costs and make manufacturing easier. Moreover, accurate calculation of complex experiments on NIF or DARHT would greatly enhance the designer's confidence in the modern codes and their description of weapons physics. The ability to respond to SFIs would also increase because most of those findings are inherently 2D or 3D in computational complexity. Their resolution is likely to be accomplished with a wider range of options than when restricted to "calibrated" weapons history.

In summary, the SSP and ASC path, laid out by DOE-NNSA and the Labs, is well thought out, has a reasonable level of program detail (particularly at the SRD level), and if followed at roughly the level envisioned in the various ASC documents, has a credible chance of achieving the milestones on approximately the predicted timescales. However, the DOE-NNSA presentations were notable for:

- The absence of very high level program representatives (however, the participation of the Head of the NNSA Office for Research and Development (R&D) and the ASC staff was exemplary);
- The lack of resource requirements needed to meet the milestones;
- The occasional view that:
 - The personnel issues of attracting and retaining people would be solved at the Labs despite declining resources and bureaucratic constraints

- Outside support from other agencies would appear because it was a good idea and was needed
- Peer review would automatically occur even in a reduced resource world, and
- That integration of ASC with other elements of the program would take place in a natural and seamless fashion.

Much more disturbing was the lack of a larger DOE-NNSA strategic plan that places future program elements in context, assigns priorities among them, and

describes the consequences of not funding various activities at a minimum critical level. In fact, ASC now plays the central integrating role previously performed by nuclear tests, and is the only arena in which all aspects of the program are tested together.

Particularly striking is the absence of mention, in testimony and other high level documents, of ASC's central role in the entire nuclear weapons program

Budget and Workforce Issues

Despite the lack of much of the detailed resource and priority information for NNSA, which is critical to our undertaking, we have seen draft budget scenarios supplied by NNSA in conjunction with the Lab planning process. We have also reviewed the history of the ASC budgets as proposed in the President's budget and then implemented in practice. This is somewhat complicated because the use of Continuing Resolutions rather than approved budgets in the recent past has made "interpretation" somewhat subjective. However, the numbers that have been made available combined with the statements in testimony and the NNSA Strategic Plan which imply scenarios of a) constant dollar future NNSA budgets at best, and b) strong priority for rebuilding the production complex provide the context for credible bounds on future ASC budgets.




Figure 1: Past and projected integrated code staff at Livermore and Los Alamos

Two charts will help illustrate the likely course of ASC funding. A history and projection of future workforce levels in integrated code efforts at Livermore and Los Alamos is shown in Figure 1 above. Integrated codes are a good proxy for the level of computational effort devoted to nuclear weapons applications. This figure demonstrates that the numbers for integrated weapons code work at the Laboratories would result in a decrease from about 170 people (at each Lab) in 2002 to approximately 70 in 2010.

At those levels, it will be very difficult to maintain the capability of many of the existing design codes, and virtually impossible to implement them on the much faster ASC computers that are planned for the next decade

A second chart is even more striking in terms of past and future ASC budgets. Fig. 2 on the following page shows the FYNSP (Fiscal Year National Security Plan) for ASC in the President's budget for FY 03 –FY14. As is evident the starting point for each fiscal year has steadily dropped from FY 03-FY09 and the five year projections in the FYNSP have little value as a predictive

tool beyond the current- and occasionally the next- fiscal year. From FY04 to FY09 the budget has declined over 25% (without including inflation), and the current FYNSP calls for another 12% drop going to FY10. If that is implemented as planned the ASC budgets will have dropped an average of more than 6% per year in a continuing slide for more than five years.

In recent years, the out-year projection of ASC budgets has been optimistic and the actual budget has often been less than the President's request




Figure 2: ASC Budget over time

At those levels viewed in terms of either manpower or dollars, it will be extremely difficult to maintain the capability of many of the existing design codes, and virtually impossible to implement them on the much faster ASC computers that are planned for the next decade. As discussed in the Computer Matters section, it will require much more effort to utilize the intrinsic power of those future machines than has been needed until now. At the projected resource levels, it is simply not achievable in a credible way, and certainly not in a fashion that retains the multiple approaches and independence needed for technical peer review. An assessment of the impact on the overall program is beyond the scope of this study, except to note that the computers and codes will not be able to reach the level of predictive capability required for significant changes to the stockpile in the future.

This Page is Intentionally Left Blank

Other DOE and National Security Missions

There are a number of areas in which DOE-NNSA uses both the capability and capacity modes of ASC machines to carry out national security mission in addition to SSP. Below are some of the most important.

Foreign Country Assessments

For many years, scientists at the weapons Laboratories have been responsible for assessing the state of the art in nuclear weapons development by other countries. In many cases, the starting point is a limited amount of intelligence information that is used to try to reverse engineer the weapons through an array of computational simulations. This is primarily done through many simplified calculations. Nonetheless, such efforts provided, and continue to provide, substantial insight into foreign country programs.

Nuclear Counter Terrorism

In many respects, nuclear counterterrorism is near the top of national security concerns because it is unlikely to be influenced by traditional deterrence or other consequences. It comes in two scenarios: the potential use of a country-built device by another group, or the construction of a primitive device or crudely assembled explosive by a group that has acquired nuclear material. Each scenario requires an extensive sequence of calculations to help in combating the threat.

The principal challenge in responding to the potential use of a country device by another group (or obviously by the country itself) is nuclear forensics in which an attempt is made to discern the explosive's origin from the debris created in the explosion. This is a very complex problem because it requires some knowledge of the details of the explosive and how those would affect the material in the vicinity of the explosion. Not only must the range of possible weapon types be simulated using the best assessments from the foreign country programs, but many calculations of possible environments (e.g. parking garages, tunnels, etc.) must also be carried out. Thus, a huge array of capacity calculations need to be done, but some capability computations are also required to validate the simpler models in complex environments.

The second problem—that of a crude device hypothetically assembled by a terrorist group—is equally daunting. The responder must try to imagine a myriad of ways in which an opponent might try to configure an explosive and then assess whether those explosives would produce a nuclear yield and/or

create radiological damage in various use scenarios. As in the example above, this requires many simple calculations that then need to be validated by a few capability computations.

In both of these situations, an equally important task is to test disablement schemes on the computer to determine whether or not one can effectively disarm an explosive without setting it off or making a radiological mess. Since the number of proposed techniques must necessarily be quite large, this also becomes a computationally intensive effort.

Vulnerabilities

A third area of interest is assessing the vulnerability of various activities and infrastructure to either a nuclear or conventional attack. For example, the Task Force heard a very comprehensive description of the issues and likely damage involved in an Electromagnetic Pulse (EMP) attack. This required a level of simulation completely impossible before ASC (and limited even now). High resolution simulations of the response of critical infrastructure ranging from bridges to nuclear power plants provide insight into how to strengthen various infrastructures and improve security (as well as improve structural robustness to guard against natural events such as earthquakes). Analogously, transporting hazardous material of various kinds can precipitate terrorist opportunities, and sequences of calculations can suggest operational or technical means of improving the safety and security of such activities (the Task Force heard several classified presentations along these lines). In all of the examples presented, the basic approach is a large array of “what if” calculations followed by detailed computations (and occasionally experiments) to verify the simpler assessments.

Work for Others

To date, most of the work described in this section has been funded at the margin by the core nuclear weapons program, with some support from the intelligence and homeland security communities. All of this funding, however, is for people using the ASC computers, and not for the machine time itself, and certainly not for the support or development of the computational infrastructure. For instance, the Purple machine at LLNL—currently used as the weapons simulation workhorse by all three Labs—is completely funded through the Defense Programs part of DOE-NNSA, and its use is administered through the weapons program at each of the Laboratories.

There are isolated instances in which there has been dedicated use of the computer capability, but these are special actions requiring approval by the head of NNSA (and not reimbursed at anywhere near full cost recovery).

Perhaps the most notable of these is the dedicated use of the Red Storm computer at Sandia to assist the U.S. Navy in shooting down an errant U.S. satellite in February 2008. For two months, the NNSA diverted Red Storm and its technical experts and codes to the classified project to simulate, assess, and plan the complex mission of shooting down the satellite. The calculations helped answer many questions including what altitude to hit the satellite, how to minimize the spread of debris (including its hazardous fuel), and the best way to ensure that the satellite was destroyed with a single shot. As with a few such instances in the past, this kind of special effort in the national interest is done without reimbursement, and all of the computations were carried out by Sandia staff. It was a heroic effort and very successful in meeting its goals.

DOE-NNSA has encouraged the Laboratories (especially Sandia) to expand their use of HPC for other national security agencies and develop a business model that provides at least partial cost reimbursement for such activities. However, the Task Force has concerns about both the functional matters associated with such potential work (who uses the computer, how the cost accounting is done so that it helps support the broad computing infrastructure, etc.) and the security questions.

In particular, a subtle security issue has to do with using the machine for different types of classified applications having different levels of classification and different objectives. While multi-level security has been a goal for a long time, it is far from being a reality. At present, the only workable policy is to require that all users of the machine are cleared for access to everything on the machine. Of course, this does not mean that they have easy access to all files, but if they should come in contact with sensitive data, damage will be limited. The question of different kinds of classifications is even trickier. While the NNSA HPC community has a deep understanding of the sensitivity of nuclear weapons-related data, only a small fraction of the technical staff have worked with intelligence-related data, which has a quite different set of sensitivities concerning the way it is handled and distributed.

The Alliance Program

Although not literally a national security effort, the Alliance program has been a very valuable and successful part of ASCI and now ASC. It consists of a set of competitive awards to university groups that apply HPC to technical problems related to weapons physics, but that are entirely unclassified. Examples include explosive astrophysical events (e.g. supernovae), turbulent flow, and simulation of accidental fires and explosions. Major research grants typically support a large computational team and center at a university for a five year period.

Both the participants and the reviewers give the Alliance program exceedingly high marks. Not only does the work meet very high scientific standards, it also has two corollary benefits for HPC in the country. First, the Lab's style of computing and large scale code development often finds its way into the academic environment and ideas from the university world also find their way back to the Labs. Both communities view this informational exchange positively. Secondly, it creates a substantial number of scientists and engineers who are now trained in the use of HPC for problem solving which is a valuable asset to our national competitiveness and ultimately to ASC and NNSA.

The Role of ASC/HPC in the Work of Other Organizations

The ASCI, and now ASC programs, have pushed the state of the art in HPC for the past 15 years (as did many of their predecessor programs during the preceding half century). Many other federal agencies and large industrial firms have benefited from this rapid advance in computing capability and exploited it for their own missions. In most cases, these other organizations use HPC technology that is a generation behind the leading edge of ASC, but in some cases they have partnerships to help pursue future computing advances.

The Task Force heard from most of the relevant federal agencies and from a number of high-end commercial users. Included below is a brief summary of the current status of HPC in those organizations and their perspective on the need for further advances.

Other Federal Agencies

DOE/ Office of Science (SC) /Office of Advanced Scientific Computing Research (ASCR) strategy is to be the leader in advancing open science through high performance computing. Their focus areas are closely aligned with DOE/SC missions: climate, bioscience, energy research, and basic science. To this end, ASCR invests broadly in HPC facilities, and in Leadership Class Facilities (LCF) at Oak Ridge National Laboratory (ORNL), Argonne National Laboratory (ANL) and the Lawrence Berkeley National Laboratory (LBNL) National Energy Research Scientific Computing Center (NERSC).

The ASCR Office also maintains well-planned, long-term investments in applied mathematics, computer science, networking, and in the DOE/SC Scientific Discovery through Advanced Computing (SciDAC) program, which includes coordinated participation by NNSA and the National Science Foundation (NSF). ASCR provides national, open computing leadership in the LCF and also the SC Innovative and Novel Computing Impact on Theory and Experiment (INCITE) allocation program. Through INCITE, the LCF currently provides extreme computing to a small number of projects selected from the general science community that have a reasonable probability of resulting in high-impact scientific discoveries. In addition, the NERSC ASCR facility provides a world-leading lower tier system that serves a much larger user community. NERSC also contributes to high-impact scientific discovery and additionally provides for the more complete exploitation of previous scientific accomplishments. Beyond simulation and computing, the mission space of NERSC includes the broad emerging HPC data-driven areas of informatics and visualization.

In the hardware arena, ASC contributes to SC in the areas of technology transfer (high end, clusters, and storage), and support of the contractor vendor base towards further development. In particular, there are active partnerships among NNSA and SC Labs including ones among Berkeley, Livermore, and Argonne with IBM and between Oak Ridge and Sandia with Cray. SC is also a member of the multi-agency High Productivity Computing Systems (HPCS) consortium. With regards to software, SC partners with ASC in joint software programs, specifically SciDAC, benefit from ASC-driven software developments that stimulate further basic and early applied research.

A recent review⁵ on the balance of activities between research and facilities by Advanced Scientific Computing Advisory Committee (ASCAC) had broad praise for the HPC activities within the SC but also recommended a greater focus on research (and software) to restore the proper balance with the efforts to develop and acquire high end facilities. As the report advised, “we must invest in facilities to stay in the game, but we must invest in research to win” (referring to our competitiveness in an international arena).

In summary, the Office of Science has taken a leadership role in developing HPC for unclassified applications in physical and life sciences. It has benefited

It is the Task Force’s view that additional leadership from the most senior levels of DOE to encourage joint efforts could enhance both the financial and technical effectiveness of the agency in pursuing both facilities and long term research and development for advanced computing

from technology derived from ASCI/ASC systems but is increasingly joining with ASC to pursue leadership class facilities. It is the Task Force’s view that additional leadership from the most senior levels of DOE to encourage joint efforts could enhance both the financial and technical effectiveness of the agency in pursuing both facilities and long term research and development for advanced computing.

DARPA began efforts to develop a new generation of economically viable, high productivity computing systems for national security and industrial user communities, following the DSB report published in 2000 on “DoD Supercomputing Needs.” The DARPA goal, to ensure U.S. lead, dominance and control in this critical technology, is enunciated in four impact areas:

1. Performance (time to solution): provide speedup critical to national security applications by a factor of 10X to 40X;
2. Programmability (idea-to-first-solution): reduce cost and time of developing application solutions;
3. Portability (transparency): insulate research and operational application software from system; and

4. Robustness (reliability): continue operating in the presence of localized hardware failure, contain the impact of software defects, and minimize likelihood of operator error.

DARPA laid out the framework for the HPCS program and is leading the effort with support from the DOE and the National Security Agency (NSA). The HPCS is currently implementing a three-phase program spanning 2002-2010. The first phase involved an industry concept study that concluded in 2003. The second phase, R&D, began in 2003 and concluded in 2006. The third phase, Development & Prototype Demo, which is scheduled to conclude in 2010, has as goal to build a petascale prototype. Funding for the petascale prototype is shared by DARPA and its consortium partners. Additionally, each vendor is contributing at least one third of the total cost of the program.

The HPCS effort is complementary to that being pursued within ASC. Its goal is to make HPC widely available at the petascale level to a broad national security community. The combined effort of ASC and HPCS will also provide important support to the computer manufacturers to continue U.S. leadership in this industry.

DoD High Performance Computing Modernization Program (HPCMP) first launched in 1992, was formalized in 1994, and began major acquisitions in 1995. Since that time, HPCMP has expanded to provide services to a wide range of DoD organizations. The HPCMP hardware strategy is to procure commercial supercomputers annually, based upon a set of quantitative and qualitative criteria, and turnover their inventory every four years. Software factors and productivity are addressed by the Productivity Enhancement and Technology Transfer program (PETs), which enables transfer of leading edge, HPC-relevant computational and computing technology onto the DoD HPCMP systems from within other parts of the DoD, and from other government, industrial and academic organizations. The DoD strategy relies on the availability of commercially available HPC machines and software which have advanced, in large measure, due to the ASC program.

Other National Security Agencies often have application sets different from those required for the NNSA-ASC mission, but the continued existence of a robust industry producing high end machines is as vital to those parts of the national security community (e.g. National Security Agency) as it is to NNSA. Machines and software produced for NNSA applications may not, themselves, be ideal for other national security problems, but the base technologies behind the ASC machines are critical for those other segments of the national security community. Direct use of ASC machines for some other national security applications could raise security issues, since the handling of different classified materials may have protocols of which not all users are aware. The Laboratories do a good job of handling these matters

within their own work but any extensions into the intelligence field will require additional measures.

NASA High End Computing (HEC) has a vision to be relied upon by NASA as an essential partner to enable rapid advances in insight and enhance mission achievements. The vision implementation strategy is to buy what is commercially available and focus on how to increase the productivity for complex systems simulation such as: modeling fluid dynamics to predict aero-thermal environments; modeling parachute deployments to examine effects of trim; and modeling Pareto-Optimal Trajectories for fuel and flight times. The NASA HEC program has two facilities with a total of four HPCs that range from 6.9 to 530 TF. To collect and share modeling expertise and experience, NASA is using a 'modeling guru' system. The guru is shared by nine communities within NASA, and allows users to work together on either wiki-based documents or binary documents and also manages documents through versioning and workflow. NASA relies on others to lead the industrial development of the top level of HPC

NSF has as strategic plan through 2010, to enable petascale science and engineering by means of deployment and support of a world-class HPC environment comprising the most capable combination of HPC assets available to the academic community. NSF is performing this through acquisition, deployment and operation of science-driven HEC systems, as well as through the development and maintenance of supporting software, new design tools, and portable, scalable applications software. NSF invests approximately \$30 million per year in hardware, funding proposals from institutions that include a vendor system with benchmarking projections. Cost-savings are achieved by leveraging pre-existing infrastructure and personnel at the 11 NSF HPC host sites throughout the U.S. As is the case for most of DoD, NSF depends primarily on other programs such as ASC to spearhead the development of top end HPC.

The Commercial Sector

The Task Force received a number of briefings that represented a wide range of views from the commercial sector, including presentations by the Council on Competitiveness, International Data Corporation (IDC) and the major industrial users.

IDC presented one of the most interesting and valuable briefings. They conducted extensive surveys of HPC users by various industrial sectors throughout the years. In particular, they surveyed the impressions HPC users, participants, vendors and other stakeholders have of ASC. Their findings indicate that ASC enjoys high marks-along the lines of "almost unprecedented in its value and execution since its creation." In short ASC is

viewed by much of the industrial world as the enterprise that has led U.S. and world computing since its inception.

The Task Force also received in-depth presentations from high end industrial users including Boeing, on aircraft design; a former Chevron executive, on oil and gas activities; Goodyear, on tire design; and, Pratt & Whitney, on engine R&D. The Task Force also received non-disclosure briefings from three of the major computer companies, including IBM, CRAY, and Intel (which are referred to in the next section of the report). In all cases, the commercial users relied completely on the government driven programs, such as ASC, to create the HPC capability that they could deploy with a lag time of a few years.

- Boeing uses advanced computing to inform and validate, in part, their aircraft designs. The aircraft industry began using computational tools in the early 1980's and has honed their skill set since. Boeing relies on HPC to make their products viable and competitive. Among the most compelling illustrations of HPC impact on Boeing's business is the significant reduction of wind tunnel tests that have now almost been entirely replaced with computational fluid dynamics modeling.
- The former Chevron executive described how major energy companies use advanced computing to support high risk exploration, as well as complex processes and associated facility designs. A key application is seismic imaging. The combination of immense datasets, low signal to noise ratios, inverse 3D propagation and many iterations make advanced computing essential. The use of HPC by major energy producers is ubiquitous and essential to their business plans. They rely heavily on the computing advances made in response to federal agency mission needs (especially NNSA-ASC) to remain competitive.
- Goodyear entered into a Cooperative Research and Development Agreement (CRADA) with SNL in 1993, in conjunction with DOE's former Tech Transfer Program in place at the time. The program enabled Goodyear to introduce a new and competitive product during a critical time of their business, and shorten their product design-to-market time from three years to a matter of months. In addition, Goodyear asserted they now save approximately \$100 million each year with product design efficiencies gained via the HPC tech transfer effort. SNL benefited from the relationship also, as they are now able to solve previously intractable weapons problems. While SNL funded the majority of the work performed under the CRADA in the first few years of the program, Goodyear began shouldering the entire cost of the partnership in 2000 which continues to be the case today. Based on their successful experience, Goodyear recommends that further consideration be made

towards continuing Laboratory tech transfer programs with industry where appropriate.

- Pratt & Whitney (P&W) finds HPC to be an essential business tool that helps them develop components and integrated systems that provide great value to their customers. More specifically, P&W has realized a reduction in development cost and schedule and an increase in their product quality through their use of HPC. While HPC is an essential enabler, there are other equally important aspects such as resolution, accuracy, speed, knowledge generation and decision-making that are important to their business. Most current computational tools are only capable of analyzing components at selected design points. A quantum jump in modeling and simulation capability is required in order to achieve the next level of capability; which would be to perform the complete component design process computationally.

Overall, ASC contributions to the applications-focused commercial sector reside centrally in areas of technology transfer, both indirectly, via hardware and associated systems and implementation software developments, and directly, through applications directed tech transfer programs such as SNL's relationship with Goodyear. The industrial base also benefits from student training as provided by the ASC Level 1 University Centers (e.g., companies like P&W seek to hire students who have hands on experience like that gained through the ASC University Centers).

Summary

The overwhelming message from all of the organizations outside of NNSA-ASC who briefed the Task Force is that ASC-developed hardware and associated software is broadly and effectively implemented. It is clear that the ASC investment is a driving engine for current U.S. HPC preeminent capability, and that impact extends far beyond the direct ASC program. The investment results in development of powerful new systems within the vendor community that see significant early application by ASC, and which are subsequently adapted several years later for considerable use by other U.S. federal agencies and commercial sector organizations.

As in the previous national security section, there is some interest in more direct use of ASC machines in a work for others context. And, as in the case of Goodyear and Sandia, there are some notable success stories in this regard. However, this can create some potential security issues; an obvious one being the question of "swinging" a machine between classified and unclassified uses. Strictly speaking, the security issue here is not the classification of the application, but rather the exposure of a previously classified machine to the open internet. While Sandia, for example, has

mechanisms for doing this that have been used on several machines for some time and has extremely careful mechanisms in place, several Task Force members have concerns that the risks associated with this strategy outweigh the accrued benefits.

This Page is Intentionally Left Blank

Computer Matters

Before delving into the detailed issues confronting the ASC program, it is useful to note some of the reasons that computer hardware and software have become such critical and difficult matters for the future of high performance computing. In the early days, until about the mid-1970s, the procurement of supercomputers involved the acquisition of a single computing machine that contained most of the important features needed for large scale computations. Some of the related functions, like reading large data bases or converting output to graphical form were done on peripheral equipment. The central computing engine became more powerful primarily by putting more transistors on a chip and arranging them in efficient ways inside the computer. Then, as now, there were only a few industrial participants: CDC, Cray, IBM and occasionally other manufacturers. Operating systems often had custom features, but the applications software tended to be rather straightforward (e.g. some version of Fortran for scientific simulation).

In the 1970s and 1980s, the vector computer represented the next step in speed and efficiency. The basic idea was that many physical systems had characteristics in which the same piece of arithmetic was performed perhaps thousands of times (e.g., calculating the stress at many points along an aircraft frame, or the wind speed at a particular height in a weather simulation). This resulted in computer architectures and software that made such operations very effective and greatly speeded up calculations that required a large number of such vector operations. Cray was particularly focused on developing computers along these lines.

As the 1980s progressed, it became clear that there were physical limitations on the number of chips that could be usefully put together to form a single computing engine, and that parallel computing—in which many small computers were connected to form the overall computer—had the greatest promise for breakthroughs in computer power and speed. An additional benefit was that the individual small computers could be sold to mass markets by the manufacturers (i.e. their normal business) and the supercomputer would only require fast communication links among the small units, not the complex design of a single custom machine with a limited number of customers. It would still need a major research and development effort on “fast interconnects” among the small computers, but this was a much simpler task than the design of a single custom HPC machine.




Figure 3: Timeline development of the fastest computers

The difficulty arose in programming applications effectively for such parallel computers. Parallel computing is optimal when each small computer can work independently of the others and is “busy” most of the time, only communicating at infrequent intervals or with a limited set of near neighbors. When greater communication and knowledge of events happening elsewhere during the simulation is necessary, then it requires very special software to take advantage of the intrinsic power of the large collection of parallel computers. Thus the list of the 500 fastest computers is increasingly a very imperfect proxy for the range of complex calculations actually done by users in different fields. It is somewhat like collecting data on an athlete’s ability to run fast and lift weights, without any regard for how well these are put together to play an actual game. Inevitably, some systems are best for some tasks than other systems and vice versa. Position on the Top 500 list is interesting and informative, but generally not determining, both because of the inadequacy of any single figure of merit and because the list encourages vendors to optimize for one benchmark.

The history of ASCI-ASC as the lead for HPC development reflects all of these factors. Vector machines are a thing of the past for ASC and the question now is how parallel will the future machines be? The “very fast” Blue Gene line of IBM computers and the hybrid Roadrunner are optimized for highly parallel applications, but are not as adept at problems requiring more frequent communication among parallel elements. Many basic science calculations are ideally suited for highly parallel work and important studies of underlying weapons science have been done on Blue Gene and Roadrunner. Conversely, the Purple machine and its envisioned successors have a smaller number of units than the Blue Gene’s but perform more effectively on weapons design calculations. For example, the Purple machine, the workhorse of weapons design, has a little over 10,000 individual computing cores with four gigabytes (GB) of memory per core, while Blue Gene has over 200,000 cores but less than a half of a GB per core. Blue Gene may win a higher place on the “fastest computer” list, but is not as easily adaptable to weapons design calculations. Nonetheless, like most lists, the evolution of “peak” computing power still has great interest, and the history of the “fastest” computer is shown in Figure 3 on the previous page. As noted earlier, nearly every one of those computers is driven by the needs of ASC and its predecessor organizations within the nuclear weapons program.

Thus, in all situations, and for a wide variety of other national security situations, the software and other custom features become extremely important in constructing a computing system that can take advantage of the intrinsically higher speed provided by Moore’s law of increasing power per chip. In developing and procuring future supercomputers it is this close relationship among parallelism, software and application sets that makes the development and procurement process very difficult and one that needs a strong iterative relationship with potential manufacturers. A much more detailed description of the evolution of the mixture of evolution and innovation in the ASCI/ASC process is given in several references^{5,6}.

The ASC Plan

When the ASCI program was established in the early 1990s, the intent and resulting requirements were based on enabling stewardship of the existing stockpile without nuclear testing. As discussed earlier in this report, this led to greatly expanded computing requirements to allow for the detailed study of nuclear phenomena in an aging stockpile.

NNSA now is required to go significantly further in two application directions: predictive capability and uncertainty quantification. We find every reason to believe that these added requirements will dramatically increase the need

for computing capability and capacity. The key questions to be addressed in this section are:

- Is the ASC plan sufficient to meet these new requirements?
- Are the likely characteristics of to-be available technology reasonably aligned with the computational requirements?

The ASC Hardware Plan

Our visits to LLNL, LANL and SNL were very helpful in providing additional detail to our understanding of DOE-NNSA efforts and more specifically, ASC activities. As a high level summary, we find the document entitled “ASC Roadmap” provides the clearest picture. As such, this section of the report will draw on that representation of the plan, specifically the discussion of Focus Area 4.

The roadmap (dated 2006) shows the transition to a National User Facility concept and focuses on computational environments for uncertainty quantification in 2007 and 2008. We found evidence of this on our field trips as the users of the computing facilities at all labs had the same point of view: the facilities were allocated based on mission needs, not the home-base of the users.

In the 2008-2012 timeframe, the roadmap shows a focus on deploying environments for weapons science studies and other capability computing needs. A 2009 target for petascale computing is included. More details were provided in a presentation at LLNL⁷. The strategy includes three categories of investments:

1. Capability Systems, which can run integrated physics codes which require large, tightly coupled architectures.
2. Capacity Systems, which allow more cost-effective computing where applications have more modest architecture requirements.
3. Advanced Architecture Systems, which explore future capability systems by increasing the risk taken and (potentially) concentrating on a subset of mission requirements.

During the 2008-2012 period, there is a dramatic change happening in the world of computing: the calculation of arithmetic operations (e.g., floating point multiply) will become dramatically cheaper while access to memory, especially, that which is large and distributed will become relatively more

expensive. If one examines the current Advanced Architecture systems (BlueGene/L and Roadrunner) as harbingers of the future this trend is readily visible. The key question is whether these advanced architecture systems will be capable of economically running the codes typically run on capability systems, or whether they will only be suitable for those applications run on capacity systems. For Uncertainty Quantification runs there is an argument to be made that this is possible. However, the same seems quite uncertain for the prediction runs.

At a high level, the current NNSA plan provides for a capability system in Fiscal Year (FY) 10 (Zia), Capacity systems in FY11 and FY14, and an Advanced Architecture system (Sequoia) in FY12. Sequoia, if successfully procured, reflects the discussion in the previous paragraph. Although designated as an advanced architecture system, it will also be aimed at capacity calculations for uncertainty quantification, and is intended to have a capability level that will be useful in many circumstances. All of these systems are under considerable pressure to reduce their capabilities and/or extend their schedules due to current budget pressure. However, if the major elements of the plan are able to be retained, there should be petaflop weapons computing available at the Laboratories within the next five years.

Further out on the roadmap, are targets for 100x petascale computing in 2016 and exascale computing in 2018. We received several presentations from potential vendors of these systems and there are several troubling trends including:

- The need for greatly increased electrical power
- The dramatic reduction in memory capacity, and
- Memory performance relative to arithmetic calculation performance

In addition to these common elements there are many differences in approaches by the different vendors, but they can't be discussed in detail here because of their proprietary nature. However, in this time frame, our concerns for the applicability of these systems for both prediction runs and uncertainty quantification are significant.

HPC Trends

Among those who track the state of high-performance computing, both nationally and internationally, there is little doubt that the NNSA investments in ASC are the largest contributor to the continuing vitality of the U.S.' current leadership in the HPC industry. At a time when other agencies lacked

the budget or the programmatic commitment to HPC, the ASC roadmap, which required a series of systems with increasing performance to meet stewardship certification milestones, and the associated procurements, ensured that multiple vendors continued to develop new systems.

One consequence of the “accelerated” part of the ASCI (now ASC) program is an emphasis on systems that could be packaged and deployed at large scale. One consequence of this is an emphasis on the commodity cluster model of HPC, to the possible detriment of alternative designs based on custom processor, interconnect and memory technology. Only via such a commodity approach could vendors deliver large-scale systems on the schedule dictated by the ASC certification milestones.

Historically, government HPC procurements have driven the very highest end of the industry. However, the dearth of purpose-built HPC designs suggests that the computing industry has shifted its focus to mid-range, commodity supercomputing, emphasizing commercial and academic markets where the majority of the users and markets lie. One indicator of this is the increasing incorporation of more commodity components. First, the community shifted from purpose-built vector systems to symmetric multiprocessors (SMPs), then to commodity clusters with custom interconnects. Today, an increasing fraction of the world’s HPC systems incorporate accelerators drawn from the computer gaming business, a true mass market. Unfortunately, these systems lack the memory and input/output bandwidth and the ease of programming needed to develop complex, multiphysics and national security applications both rapidly and efficiently. This is a worrisome trend that does not bode well for the future of national security needs.

Site Issues

There appears to be some concern about physical location and management of machines. In our view, location should be a non-issue. It is not necessary or even desirable to have users close to the machine and machines at every site. Obviously building and maintenance costs increase and additional machines may be viewed by decision-makers as instances of replication.

Management of machines across sites is a more complex but not intractable question. In fact, our discussions with users at the three sites show that the machines are already managed as complex-wide resources with users from all labs running at least part of their workload on all resources. The LANL/SNL cooperation focused on the Zia system is very strong evidence that cooperation between the labs can be deeper and more extensive than it

already is. If, as expected, budgets decline, such cooperation will be essential for the Labs to continue to succeed in the core ASC mission.

Impact of ASC Investment on Vendor Plans

As discussed previously in this section, and confirmed by our discussions with potential vendors, future development of leading edge computer capabilities will be increasingly difficult and will require strong interactions between potential customers and the computer companies. It will be the specific national security or scientific applications that will often dictate particular technology choices. There is no generic next generation of HPC that can be developed without tight coupling to mission needs. Consequently, ASC investment, with support and partnerships from the DOE Office of Science and DARPA's HPCS program, will largely dictate the high end computer development and competitiveness in the U.S.

For most other federal agencies and commercial users, the acquisition of a capability-level ASC machine cannot be justified within their mission and resource constraints. First, the capital expense is large and at a scale where

Unless conditions change, ASC is in danger of becoming an unattractive place for the best and brightest computer scientists to practice their profession.

normal prudence requires fairly sure returns on the investment. Secondly, there are few problems where the increased performance of a capability over a capacity machine can compensate for the substantially increased costs. Industrial and many federal agency users will tend to be satisfied with machines a


generation or so behind the highest end as these machines are often cheaper and have the programming difficulties already worked out by those whose missions require a capability level of performance.

Workforce Issues

Since inception, the U.S. weapons design workforce has consisted of dedicated individuals with world-class, advanced scientific or engineering education, many of whom have committed their entire professional careers to the enterprise. The highly specialized nature of nuclear weapons design requires that there be continuity in the workforce and critical mass in its size. Moreover, the workforce must be drawn from the smaller pool of talented U.S. citizens who can be cleared for the sensitive nature of the work⁸.

Several factors may adversely affect the ability of the computation-centric weapons design program to sustain the required size and continuity of the workforce. Over the past 50 years, the overall proportion of science,

FIGURE 3-6. Non-U.S. citizens' share of doctorates awarded, by major field: 1960–64 and 1995–99



SOURCE: NSF/NIH/USED/NEH/USDA/NASA, Survey of Earned Doctorates and Doctorate Records File.

Figure 4: The graph above can be found at <http://www.nsf.gov/statistics/nsf06319/chap3.cfm#sect8>

engineering, mathematics and computer science Ph.D.'s conferred to U.S. citizens has decreased significantly (see Figure 4 above). Over the same time span, U.S. industry and academia demands for advanced degrees in areas competing with weapons design have increased significantly, creating more competitive pressure on hiring.

Current ASC computing capabilities rely critically on computer science disciplines such as high performance parallel computing, scientific computing, specialized compiler techniques and large-scale data storage and networking technology. In addition scientific computing expertise in physics, materials science and mechanical engineering is of particular relevance.

Laboratory staffing trends for computational science at the weapons Laboratories are troubling. As noted in Figure 1 on page 16, the staff levels for code development have dropped by nearly two-thirds in less than a decade. There is considerable anecdotal evidence of a flow of talented computational scientists to the Office of Science labs which are now joining the forefront of computing. In part this is due to the diminishing resources at the NNSA labs, and in some measure, because of the security and bureaucratic restriction constraints at the weapons labs. The recruitment of first class computer professionals is a highly competitive activity, and the relatively low salaries in government labs can only be offset by the opportunity to work on really challenging and important problems with

adequate resources. Unless these conditions are addressed by the broader program, ASC is in danger of becoming an unattractive place for the best and brightest computer scientists to practice their profession.

This Page is Intentionally Left Blank

References

1. Mara, G.L. and Goodwin, B.T. "Stewarding a Reduced Stockpile," LLNL-CONF-403041, 2008.
2. Kusnezov, D.F. and Frazier, N. "Advanced Simulation and Computing ROADMAP, NA-ASC-105R-06-Vol.1, 2006.
3. Meisner, R. "A Platform Strategy for the Advanced Simulation and Computing Program", NA-ASC-113R-07-Vol.1, 2007.
4. Kusnezov, D.F., "Advanced Simulation and Computing: the Next Ten Years", NA-ASC-100R-04, 2004.
5. International Competitiveness: Facilities for Participation Research, for Leadership, Report of the ASCAC Balance Panel, 2008
6. McCoy, M. "Riding the Waves of Supercomputing Technology", internal LLNL paper, 2003.
7. Seager, M. "The ASC Sequoia Programming Model", internal LLNL paper, 2008.
8. Chiles, H.G. (chair), DSB Task Force Report on Nuclear Deterrence Skills, 2008. (See this report for more information of competitive hiring pressures)

This Page is Intentionally Left Blank

Appendix A: Terms of Reference

MEMORANDUM FOR THE CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: Terms of Reference – Defense Science Board (DSB) Task Force on the
National Nuclear Security Administration (NNSA) Strategic Plan for Advanced
Computing

The DSB shall conduct an evaluation of the strategic plan for advanced computing of the NNSA and assess the impact of using the planned capability for other National Security issues.

Advanced computing capabilities have long supported a wide range of National Defense issues. In addition to traditional uses, advances in supercomputing capability are creating exciting new opportunities in basic scientific research that can be employed to generate breakthroughs for national security applications. Furthermore, the increasing complexity and the large span of issues is driving the development of validated tools that help address and solve time-urgent issues by developing new and/or unique national security mission capabilities. These validated tools can only be derived from advanced computing capabilities.

The NNSA employs advanced computing capabilities for a specific capability. Under the auspices of the Advanced Simulation and Computing (ASC) program, NNSA addresses nuclear weapons stockpile and national security issues through the development and use of computer simulations. The ASC integrated codes incorporate high-fidelity scientific models validated against experimental results and compared to theory. The mission of the ASC program is to fulfill the science-based simulation requirements of the Stockpile Stewardship Program, which underpins NNSA efforts to certify the safety, performance, and reliability of nuclear weapons.

Problems include: advanced design and manufacturing processes; understanding accident scenarios; nuclear weapons aging; and the resolution of concerns that arise when older weapons are opened up and inspected. The complexity of the physics associated with nuclear weapons science has driven the need for state-of-the-art computer capabilities. Consequently, for the past half-century, some of the world's fastest computers typically resided at the weapons laboratories.

The Task Force should conduct an evaluation that shall include the following:


- (1) An assessment of:
 - (A) the adequacy of the strategic plan in supporting the Stockpile Stewardship Program;

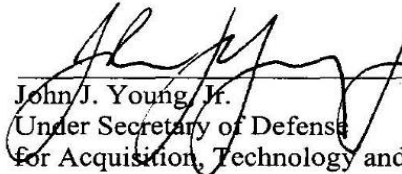
- (B) the role of research into, and development of, high-performance computing supported by the NNSA in fulfilling the mission of the NNSA and in maintaining the leadership of the United States in high-performance computing;
 - (C) the impacts of changes in investment levels or research and development strategies on fulfilling the missions of the NNSA; and
 - (D) the importance of the NNSA and partner agencies using current and projected scientific computing capabilities to address a broad spectrum of national security challenges, including threats to citizens and to the Nation's infrastructure.
- (2) An assessment of the efforts of the Department of Energy to:
- (A) coordinate high-performance computing work within the Department of Energy, in particular between the NNSA and the Office of Science;
 - (B) develop joint strategies with other Federal agencies and private industry groups for the development of high-performance computing; and
 - (C) share high-performance computing developments with private industry and capitalize on innovations in private industry in high-performance computing.

The Task Force shall have access to all levels of classified information needed to develop its assessment and recommendations. A report shall be submitted to the Secretary of Energy and Secretary of Defense with sufficient lead time to meet the legislative deadline for the report to Congress.

The Study will be sponsored by me as the Under Secretary of Defense for Acquisition, Technology and Logistics; the Administrator, National Nuclear Security Administration; and the Acting Assistant to the Secretary for Nuclear, Chemical and Biological Programs. Mr. Bob Nesbit and Dr. Bruce Tartar will serve as the Task Force co-Chairmen. Ms. Jacqueline Bell, Defense Threat Reduction Agency, and Dr. Dimitri Kusnezov, NNSA, will serve as the co-Executive Secretaries. Major Charles Lominac, USAF, will serve as the DSB Military Assistant.

The Task Force will operate in accordance with the provisions of P.L. 92-463, the "Federal Advisory Committee Act," and DoD Directive 5105.4, the "DoD Federal Advisory Committee Management Program." It is not anticipated that this Task Force will need to go into any "particular matters" within the meaning of title 18, United States Code, Section 208, nor will it cause any member to be placed in the position of action as a procurement official.

 Date 4/3/08
Thomas P. D'Agostino
Administrator
National Nuclear Security Administration

 Date 4/8/08
John J. Young, Jr.
Under Secretary of Defense
for Acquisition, Technology and Logistics

This Page is Intentionally Left Blank

Appendix B: Task Force Membership

Co-CHAIRMEN

| | |
|-------------------|---|
| MR. ROBERT NESBIT | The MITRE Corporation |
| DR. BRUCE TARTER | Lawrence Livermore National Laboratory (Ret.) |

EXECUTIVE SECRETARIATS

| | |
|----------------------|---|
| DR. DIMITRI KUSNEZOV | NATIONAL NUCLEAR SECURITY ADMINISTRATION |
| MS. JACQUELINE BELL | DEFENSE THREAT REDUCTION AGENCY |

MEMBERS

| | |
|----------------------|---|
| DR. JOHN BOISSEAU | UNIVERSITY OF AUSTIN, TEXAS |
| DR. WILLIAM CARLSON | IDA CENTER FOR COMPUTING SCIENCES |
| DR. GEORGE CYBENKO | DARTMOUTH COLLEGE |
| DR. JILL DAHLBURG | NAVAL RESEARCH LABORATORY |
| DR. SIDNEY KARIN | UNIVERSITY OF CALIFORNIA, SAN DIEGO (RET.) |
| DR. DANIEL REED | MICROSOFT |
| DR. FRANCIS SULLIVAN | IDA CENTER FOR COMPUTING SCIENCES |
| DR. VALERIE TAYLOR | TEXAS A&M |
| DR. PETER WEINBERGER | GOOGLE |

GOVERNMENT ADVISORS

| | |
|--------------------|---|
| MR. ROBERT MEISNER | NATIONAL NUCLEAR SECURITY ADMINISTRATION |
| DR. KAREN PAO | NATIONAL NUCLEAR SECURITY ADMINISTRATION |

DEFENSE SCIENCE BOARD

| | |
|--------------------------------|-----------------------|
| Mr. Brian Hughes | Defense Science Board |
| Lt. Col. Charles Lominac, USAF | Defense Science Board |

SUPPORT

| | |
|---------------------|---|
| MS. MICHELLE ASHLEY | SCIENCE APPLICATIONS INTERNATIONAL CORPORATION |
| MR. CHRIS GRISAFE | SCIENCE APPLICATIONS INTERNATIONAL CORPORATION |
| MS. AMELY MOORE | SCIENCE APPLICATIONS INTERNATIONAL CORPORATION |
| MS. LAUREN YORK | SCIENCE APPLICATIONS INTERNATIONAL CORPORATION |

This Page is Intentionally Left Blank

Appendix C: List of Briefings

April 16-17, 2008

| Organization | Title |
|--|--|
| Defense Science Board | DSB Administrative Brief |
| National Nuclear Security Administration | Program Overview |
| Lawrence Livermore National Laboratory Los Alamos National Laboratory | Nuclear Weapons Certification & Assessment |
| National Nuclear Security Administration Sandia National Laboratory Lawrence Livermore National Laboratory Los Alamos National Laboratory | National Security Applications |
| University of Illinois at Urbana-Champaign | National Academy 2005 Study "The Future of Supercomputing" |
| University of Texas | NSF 2006 Study on Simulation Based Engineering Sciences |
| National Nuclear Security Administration Sandia National Laboratory Lawrence Livermore National Laboratory | Collaborations, Partnerships and Investment Strategies |

June 22-23, 2008

| | |
|----------------------------|-------------------------------|
| IBM Cray, Inc. Intel | Industry Roadmaps to Exaflops |
|----------------------------|-------------------------------|

July 30-31, 2008

| | |
|--|---|
| Department of Defense | Ethics Briefing |
| Defense Advanced Research Projects Agency National Security Agency Department of Defense | Roadmaps and Strategic Planning for DoD |
| National Science Foundation Advanced Scientific Computing Research | Strategic and Program Plans for DOE and NSF |
| Stanford University University of Utah | Academic Alliances |
| Lawrence Livermore National Laboratory | Stockpile Stewardship at LLNL Overview |
| Lawrence Livermore National Laboratory | Boost |
| Lawrence Livermore National Laboratory | Energy Balance |
| Lawrence Livermore National Laboratory | Secondary Performance |
| Lawrence Livermore National Laboratory | Uncertainty Quantification (UQ) Requirements |
| Lawrence Livermore National Laboratory | Multiscale Modeling in Support of Weapons |
| Lawrence Livermore National Laboratory | National User Facility; the Purple Capability Computing Campaigns |
| Lawrence Livermore National Laboratory | Sequoia Procurement |
| Lawrence Livermore National Laboratory | Terascale Simulation Facility (TSF) Tour |
| Lawrence Livermore National Laboratory | Underground Facility Defeat |
| Lawrence Livermore National Laboratory | Traumatic Brain Injury |
| Lawrence Livermore National Laboratory | Bioinformatics |
| Lawrence Livermore National Laboratory | National Ignition Facility (NIF) Tour |

| | |
|--|---|
| Lawrence Livermore National Laboratory | Outputs, Electro Magnetic Pulse (EMP) and Effects |
| Lawrence Livermore National Laboratory | Nuclear Forensics |
| Lawrence Livermore National Laboratory | Institutional Computing |
| Lawrence Livermore National Laboratory | Energy Security |

August 18-19, 2008

| | |
|--------------------------------|---|
| Los Alamos National Laboratory | Welcome |
| Los Alamos National Laboratory | Future Directions for Stewardship Computing and Simulation |
| Los Alamos National Laboratory | Roadrunner and the Future of Applications Programming |
| Los Alamos National Laboratory | Energy Balance Simulation Studies |
| Los Alamos National Laboratory | 3D Boost Simulation Studies |
| Los Alamos National Laboratory | Capability Computing and the SNL ACES Partnership |
| Los Alamos National Laboratory | Atomistic Simulations for Predictability; MD Ejecta Studies |
| Los Alamos National Laboratory | Application of LANL Nuclear Weapons Capabilities to Nuclear Counter-Terrorism & Intelligence Programs |
| Los Alamos National Laboratory | Urban Explosion Consequence Assessment |
| Los Alamos National Laboratory | Urban Nuclear Consequence Management |
| Sandia National Laboratory | Welcome and Review of the Day's Agenda |
| Sandia National Laboratory | SNL ASC Overview and DSW Alignment |

| | |
|----------------------------|--|
| Sandia National Laboratory | Survivability |
| Sandia National Laboratory | QASPR |
| Sandia National Laboratory | Special Application |
| Sandia National Laboratory | Tour of Red Storm and Discussion of Operations |
| Sandia National Laboratory | When Life Deals You Lemons |
| Sandia National Laboratory | Safety |
| Sandia National Laboratory | OPUS |
| Sandia National Laboratory | Electromagnetic Applications |
| Sandia National Laboratory | ZR Applications |

October 2-3, 2008

| | |
|--|--|
| National Nuclear Security Administration | Final Remarks to DSB Study on NNSA Supercomputing |
| NASA | NASA's Computational Modeling Challenges |
| IDC | A Study of the ASC Program's Effectiveness in Stimulating HPC Innovation |
| Council on Competitiveness | Industrial Applications |
| Energy and Technology Strategies | Barrels and Bytes: Industrial Computing for Oil and Gas |
| Pratt & Whitney | A Perspective From Gas Turbine Industry |
| Boeing | Will we ever run out of the need for more detailed calculation? |
| Good Year | Analysis-Based Design: The Goodyear Story |
| Google | Inside the Cloud |

Appendix D: Acronyms and Initialisms

A

ANL: Argonne National Laboratory

ASCI: Accelerated Strategic Computing Initiative

ASC: Advanced Simulation and Computing

ASCR: Advanced Scientific Computing Research

C

CRADA: Cooperative Research and Development Agreement

CTBT: Comprehensive Test Ban Treaty

D

DARHT: Dual Axis Radiographic Hydrodynamic Test Facility

DoD: Department of Defense

DARPA: Defense Advanced Research Projects Agency

DOE: Department of Energy

E

ENIAC: Electronic Numerical Integrator And Computer

Exeflop: 10^{18} Floating Operations Per Second (See FLOPS)

EMP: Electromagnetic Pulse

F

FY: Fiscal Year

FLOPS: Floating point Operations Per Second is a measure of a computer's performance, especially in fields of scientific calculations that make heavy use of floating point calculations, similar to instructions per second.

H

HEC: High End Computing

HPCMP: High Performance Computing Modernization Program

HPC: High Performance Computing

HPCS: High Productivity Computing Systems

L

LANL: Los Alamos National Laboratory

LCF: Leadership Class Facilities

LEP: Life Extension Program

LBNL: Lawrence Berkeley National Laboratory

LLNL: Lawrence Livermore National Laboratory

N

NASA: National Aeronautics Space Administration

NERSC: National Energy Research Scientific Computing Center

NSA: National Security Agency

NTS: Nevada Test Site

NNSA: National Nuclear Security Administration

NIF: National Ignition Facility

NSF: National Science Foundation

O

ORNL: Oak Ridge National Laboratory

P

PCF: Predictive Capability Framework

P&W: Pratt and Whitney

Q

QMU: Quantification of Margins and
Uncertainties

R

R&D: Research and Development

SLBM: Submarine Launched Ballistic
Missile

S

SMPs: Symmetric Multiprocessors

SSP: Stockpile Stewardship Program

SFIs: Significant Findings

V

V&V: Validation and verification

Other

2D: Two Dimensional

PCF: Predictive Capability Framework

Petaflop: 10^{15} FLOPS

RRW: Reliable Replacement Warhead

SNL: Sandia National Laboratory

SRD: Secret Restricted Data

SC: Office of Science

3D: Three Dimensional

This Page is Intentionally Left Blank



*Office of the Under Secretary of
Defense for Acquisition, Technology and Logistics
Washington, D.C. 20301-3140*